BBC RD 1975/31

RESEARCH DEPARTMENT

REPORT

PULSE CODE MODULATION OF VIDEO SIGNALS: visibility of level quantising effects in processing channels

M. Weston, B.A.

Research Department, Engineering Division
THE BRITISH BROADCASTING CORPORATION October 1975

BBC RD 1975/31
UDC 621.397
621.376.56

PULSE CODE MODULATION OF VIDEO SIGNALS: VISIBILITY OF LEVEL QUANTISING EFFECTS IN PROCESSING CHANNELS

M. Weston, B.A.

Summary

Previous work has shown the number of bits per sample needed to represent gamma-corrected television signals. A theoretical analysis extends these results to non-gamma-corrected and logarithmic signals which occur in signal origination equipment. This analysis shows that level quantisation of non-gamma-corrected signals is most visible in dark areas of the picture, the degree of visibility depending upon the viewing conditions. Experiments confirm these results and show that 8 bits with dither give only just acceptable pictures, and that 10 bits per sample may be required to provide a performance that is satisfactory in all circumstances.

Issued under the authority of

P. Rainger

Head of Research Department

Research Department, Engineering Division,
BRITISH BROADCASTING CORPORATION

October 1975
(PH-148)

PULSE CODE MODULATION OF VIDEO SIGNALS: VISIBILITY OF LEVEL QUANTISING EFFECTS IN PROCESSING CHANNELS

PULSE CODE MODULATION OF VIDEO SIGNALS: VISIBILITY OF LEVEL QUANTISING EFFECTS IN PROCESSING CHANNELS

M. Weston, B.A.

1. Introduction

Processing channels pre-correct for the non-linearity (gamma-law) of television displays by applying to the television signal an approximation to the inverse non-linearity (gamma correction). Such channels also matrix the colour signals to improve colour reproduction and boost certain components of the signal in order to improve resolution (aperture correction).

Modern analogue channels perform these functions adequately but it is expected that a digital channel would have the following advantages:

(i) Improved accuracy of processing.
(ii) Improved stability, reducing routine adjustment.
(iii) Improved reliability.
(iv) Reduced noise and spurious signals (especially in the vertical aperture corrector).
(v) Easier to interface with other digital processing.
(vi) Easier to manufacture and commission.
(vii) Possible future cost savings.
(viii) Easier to elaborate so as to provide more advanced processing without sacrificing stability.

In a digital television system the brightness of each point in the picture is represented by a binary number, thus restricting the originally continuous range of brightness to a limited number of brightness levels. This is known as quantisation.

If the difference in brightness between levels is too great, sharp lines (known as contours) are visible where the brightness changes from one level to the next.

Previous work has indicated the number of evenly spaced levels required to conceal the effects of the quantisation of gamma-corrected signals.[1] This Report determines both theoretically and practically the number of levels needed to represent signals occurring in various parts of a processing channel.

2. Theory – monochrome

2.1. Visibility

The visibility of a luminance change (such as that occurring at a contour) depends on the size of the change ΔY and on the mean luminance Y at which the change occurs, and it has been shown that the transition will be invisible if the fractional change $\Delta Y/Y$ in 'perceived' luminance Y is less than about 2%. In this Report $\Delta Y/Y$ (fractional change in 'perceived' luminance) is used to indicate the visibility of a contour.

$\Delta Y/Y$ is a good measure of visibility, over the contrast range of television, only if Y is the 'perceived' luminance: i.e. the luminance perceived at the retina of the observer's eye. Unfortunately, this is not an easily measured quantity; it depends not only on the electrical signal applied to the display but also on the amount of stray stimulation of the eye. This stray stimulation comes from imperfections of the display (e.g. flare and reflection of ambient light from the display face) and from imperfections of the eye itself (e.g. flare, lag and background neural activity*). These effects lead to an apparent decrease in sensitivity in dark areas unless they are taken into account when calculating $\Delta Y/Y$.

This Report combines all these effects together as a single stray stimulation term F which is added to the luminance Y_d of an 'ideal' display to give the 'perceived' luminance $Y = Y_d + F$.

The value of F at any point in the picture depends on the nature of the rest of the picture (because of flare), on the previous picture (because of lag) and on the ambient illumination (which is reflected from the display). F may thus vary over a fairly wide range. It is therefore wise to perform subjective tests whenever theory predicts that the overall visibility of quantisation depends on F. The results of tests by Geddes[2] using patterns somewhat similar to those caused by quantisation superimposed on test card C can be explained by values of F around 2–4%. The theoretical graphs in this Report are all plotted for values of F of 0%, 2% and 4%.

2.2. Quantisation of display luminance Y_d

Display luminance Y_d is the luminance of an 'ideal' display; it is not an electrical signal. However, it is proportional to the input of an 'ideal' gamma corrector (which exactly compensates the gamma law of the ideal display). A detailed analysis of input signal quantisation is given in Section 2.5; however, a simplified outline of the effect of quantisating the input to a processing channel is given below.

If the luminance Y_d of the displayed picture is quantised into a set of evenly-spaced levels the difference between successive levels $\Delta Y_d = (Y_{d\ max})/(2^n)$ (where n is the number of bits per sample used to code the luminance range from zero to $Y_{d\ max}$).

If the eye and the display were free of flare the perceived luminance Y would be identical to the displayed luminance Y_d and the fractional change in perceived luminance would be

* This may not be an exact model of the eye but it explains most of the observed effects reasonably accurately.

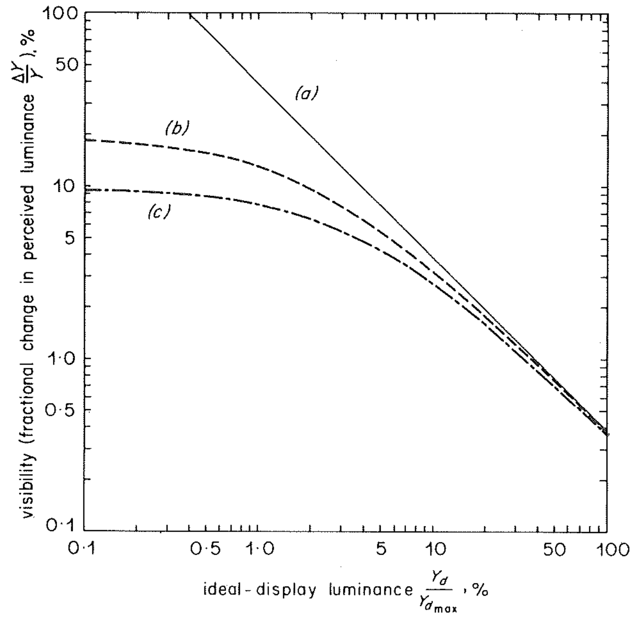

Fig. 1 - Visibility of 8 bit luminance quantisation
(*a*) ———— $F = 0$ (*b*) - - - - - - $F = 2\%$ (*c*) — — — $F = 4\%$

$$\frac{\Delta Y}{Y} = \frac{\Delta Y_d}{Y_d} = \frac{1}{2^n} \cdot \frac{Y_{d\,max}}{Y_d}$$

Thus $\Delta Y/Y$, which is a measure of the visibility of the transition ΔY, would be inversely proportional to the display luminance Y_d. This relationship is shown (for n = 8) in Fig. 1, curve (*a*), where $\Delta Y/Y$ is plotted against $(Y_d/Y_{d\,max})$ on log axes.

It will be seen that $\Delta Y/Y$ tends to infinity as Y_d tends to zero, so that however many bits were used in this idealised case quantisation would always be visible if the darkest parts of the picture were sufficiently dark.

In practice, the perceived luminance Y is not identical to the displayed luminance Y_d since the eye is subject to the various forms of stray stimulation, discussed in Section 2.1.

Under these conditions, $\dfrac{\Delta Y}{Y} = \dfrac{\Delta Y_d}{Y_d + F} = \dfrac{1}{2^n} \cdot \dfrac{Y_{d\,max}}{(Y_d + F)}$

Examples of this relationship are shown in Fig. 1, curves (*b*) and (*c*), for values of F of 2% and 4% of $Y_{d\,max}$ respectively.

The maximum value of $\Delta Y/Y$ occurs when $Y_d = 0$ and is equal to:

$$\frac{\Delta Y}{Y} = \frac{1}{2^n}\left(\frac{Y_{d\,max}}{F}\right)$$

Thus if the luminance (or ideal gamma-corrector input) is quantised into evenly spaced levels, contouring is most visible in the darkest parts of the picture. The visibility depends on the value of F (i.e. on the stray stimulation of the eye).

The precise value of F is not known since it depends on the properties of the eye, the viewing conditions and the picture displayed. Subjective tests were therefore performed to find the number of bits required in practice. These are described in Sections 4 and 5.

2.3. Quantisation of gamma-corrected signals $V = Y_d^{1/\gamma}$

The change ΔY_d in ideal display-luminance Y_d produced by a small change ΔV in the gamma-corrected signal V can be determined by differentiation:

Since $Y_d = V^\gamma$ (where γ is the gamma of the display)

$$\delta Y_d = \frac{\gamma V^\gamma}{V} \cdot \delta V$$

$$= \gamma Y_d \frac{\delta V}{V}$$

$\therefore$ For a small change ΔV (e.g. that due to quantising)

$$\Delta Y_d \simeq \gamma Y_d \cdot \frac{\Delta V}{V}$$

But since $Y = Y_d + F$ and $\Delta Y = \Delta Y_d$ (F is constant, so ΔF is ignored)

$$\frac{\Delta Y}{Y} = \frac{\Delta Y_d}{(Y_d + F)} = \frac{\gamma}{\left(1 + \dfrac{F}{Y_d}\right)} \cdot \frac{\Delta V}{V}$$

If the gamma-corrected signal V is quantised into a set of evenly spaced levels the difference between successive levels $\Delta V = (V_{max})/(2^n)$.

$$\therefore \quad \frac{\Delta Y}{Y} = \frac{\gamma}{2^n} \cdot \frac{V_{max}}{V} \cdot \frac{1}{\left(1 + \dfrac{F}{Y_d}\right)}$$

$$\frac{\Delta Y}{Y} = \frac{\gamma}{2^n} \cdot \left(\frac{Y_{d\,max}}{Y_d}\right)^{1/\gamma} \cdot \frac{1}{\left(1 + \dfrac{F}{Y_d}\right)}$$

(Note that if $\gamma = 1$ this gives the same result as in Section 2.2).

This relationship is shown (for γ = 2·8 and n = 8) in

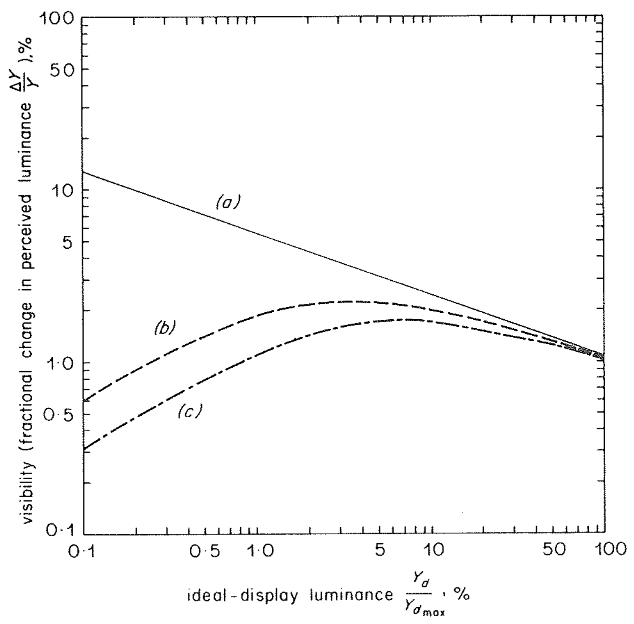

Fig. 2 - Visibility of 8 bit gamma-corrected quantisation
(a) ——— $F = 0$ (b) – – – – $F = 2\%$ (c) —·—· $F = 4\%$

Fig. 2, curves (*a*), (*b*) and (*c*) for values of F of 0%, 2% and 4% of $Y_{d\,max}$ respectively.

It can be shown that the maximum value of $\Delta Y/Y$ occurs when $Y_d = (\gamma - 1)F$ and is:

$$\left(\frac{\Delta Y}{Y}\right)_{max} = \frac{(\gamma-1)^{(\gamma-1)}}{2^n}\left(\frac{Y_{d\,max}}{F}\right)^{1/\gamma}$$

e.g. if $\gamma = 2{\cdot}8$

$$\left(\frac{\Delta Y}{Y}\right)_{max} = \frac{1{\cdot}46}{2^n}\left(\frac{Y_{d\,max}}{F}\right)^{1/2{\cdot}8}$$

Thus the maximum visibility is less dependent on stray stimulation of the eye (F) than is the case for quantisation of Y_d.

Curves (*b*) and (*c*) of Fig. 2 show that the maximum visibility of the effect of 8 bit quantisation of gamma-corrected signals is approximately 2% (i.e. just perceptible). This agrees with earlier tests by V.G. Devereux,[1] who showed that the effects of 9 bit quantisation of composite signals (i.e. including syncs and subcarrier) are just perceptible.

2.4. Quantisation of logarithmic signals $L = \log_e Y_d$

It is common practice in telecine signal processing to perform masking and gamma-correction by matrixing and attenuating logarithmic signals. Digital processing channels may also use logarithmic processing, so it is necessary to know how many bits are required to represent the logarithmic signals.

The change ΔY_d in ideal display-luminance Y_d produced by a small change ΔL in a logarithmic signal L can be determined by differentiation:

Since $L = \log_e Y_d$

$$Y_d = e^L$$

$$\delta Y_d = e^L \delta L = Y_d \delta L$$

$$\therefore \quad \Delta Y_d = Y_d \Delta L$$

But since $Y = Y_d + F$ and $\Delta Y = \Delta Y_d$

$$\frac{\Delta Y}{Y} = \frac{Y_d}{Y_d + F}\Delta L = \frac{1}{\left(1 + \dfrac{F}{Y_d}\right)} \, . \, \Delta L$$

If the logarithmic signals are quantised over a range from $\log_e(Y_{d\,min})$ to $\log_e(Y_{d\,max})$ the difference between successive levels $\Delta L = (1)/(2^n) \log_e(Y_{d\,max}/Y_{d\,min})$

$$\therefore \quad \frac{\Delta Y}{Y} = \frac{1}{2^n} \, . \, \log_e\left(\frac{Y_{d\,max}}{Y_{d\,min}}\right) \, . \, \frac{1}{\left(1 + \dfrac{F}{Y_d}\right)}$$

$\Delta Y/Y$ is plotted against $(Y_d/Y_{d\,max})$ (for $n = 8$ and $(Y_{d\,max}/Y_{d\,min}) = 1000$) in Fig. 3, curves (*a*), (*b*) and (*c*) for values of F of 0%, 2% and 4% of $Y_{d\,max}$. The maximum value of $\Delta Y/Y$ occurs when Y_d is a maximum and is approximately

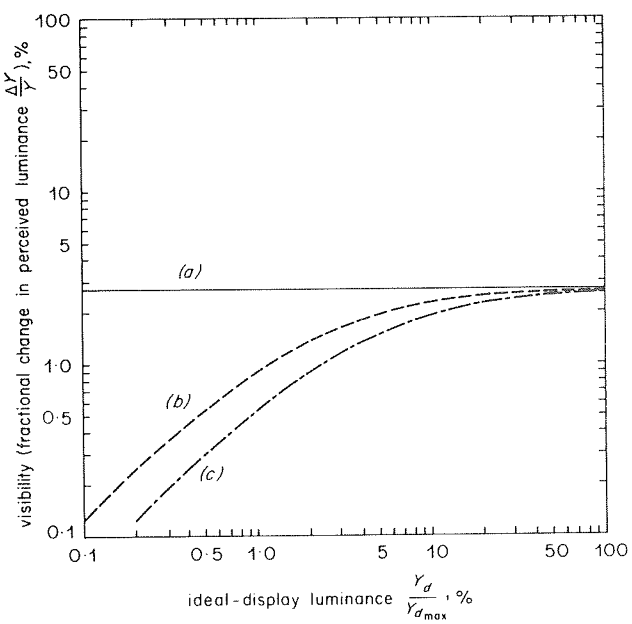

Fig. 3 - Visibility of 8 bit logarithmic quantisation
(a) ——— $F = 0$ (b) – – – – $F = 2\%$ (c) —·—· $F = 4\%$

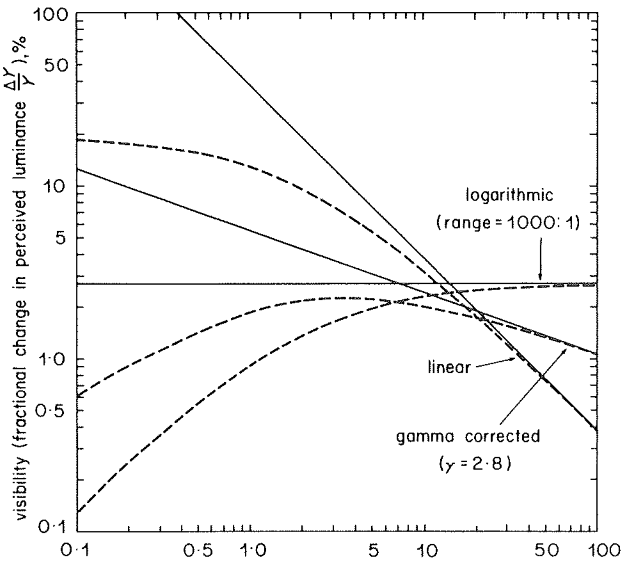

Fig. 4 - Visibility of 8 bit quantisation
(a) ——— $F = 0$ (b) - - - - $F = 2\%$

$$\left(\frac{\Delta Y}{Y}\right)_{max} = \frac{1}{2^n} \log_e \left(\frac{Y_{d\,max}}{Y_{d\,min}}\right)$$

which is independent of F.

2.5. Quantisation of the channel input $P = (Y_d^{\;1/\gamma} - s)^{\lambda}$

Fig. 4 summarises the results so far obtained for 8 bit quantisation. It will be seen that quantisation of logarithmic signals (over a range of 1000:1) has approximately the same maximum visibility as quantisation of gamma-corrected signals for ($F = 2\%$) and therefore broadly requires the same number of bits to make the effects of quantisation imperceptible.

By comparison, the effects of the quantisation of signals that are proportional to the 'linear' display luminance are much more visible. It is thus the coding of these signals which poses the greatest problem, since a large number of bits is required if quantisation is not to be visible in the darkest parts of the picture. Such signals are found at the input to a camera or telecine gamma-corrector.

Fig. 5 shows a basic signal chain, starting with the input from the photosensitive device and ending with the viewer's eye. A black-level (sit) control has been included since the adjustment of this control can effect the maximum visibility of input quantisation. (Also the gamma-corrector power law, $1/\lambda$, is not necessarily the precise inverse of γ.)

As before, the visibility of small changes ΔP in the input signal P can be determined by differentiation.

From Fig. 5, $Y_d = (P^{1/\lambda} + s)^{\gamma}$

Differentiating gives:

$$\frac{dY_d}{Y_d} = \frac{\gamma}{\lambda} \cdot \frac{P^{1/\lambda}}{(P^{1/\lambda} + s)} \cdot \frac{dP}{P}$$

since $Y = Y_d + F$ and $\Delta Y = \Delta Y_d$

$$\frac{\Delta Y}{Y} = \frac{1}{\left(1 + \frac{F}{Y_d}\right)} \cdot \frac{\Delta Y_d}{Y_d}$$

$$= \frac{1}{1 + \frac{F}{(P^{1/\lambda} + s)^{\gamma}}} \cdot \frac{\gamma}{\lambda} \cdot \frac{P^{1/\lambda}}{(P^{1/\lambda} + s)} \cdot \frac{\Delta P}{P}$$

$$= \frac{\gamma}{\lambda} \cdot \frac{(P^{1/\lambda} + s)^{(\gamma - 1)}}{\left((P^{1/\lambda} + s)^{\gamma} + F\right)} \cdot \frac{\Delta P}{P^{(1 - 1/\lambda)}}$$

If the input signal P is quantised using 2^n evenly spaced levels the difference between successive levels $\Delta P = (P_{max})/(2^n)$

$$\frac{\Delta Y}{Y} = \frac{\gamma}{\lambda} \cdot \frac{(P^{1/\lambda} + s)^{(\gamma - 1)}}{(P^{1/\lambda} + s)^{\gamma} + F} \cdot \frac{P_{max}}{P^{(1 - 1/\lambda)}} \cdot \frac{1}{2^n}$$

This is plotted in Fig. 6 for $\lambda = \gamma = 2{\cdot}8$; $F = 2\%$ of P_{max}; n = 8 and various values of s. It will be seen that if s is zero the curve is identical to Fig. 1, curve (b) (as would be expected, since in this case $Y_d = P$). If, however, s is positive the visibility of quantisation effects in dark areas is enhanced since the large steps produced by the black-stretching action of the gamma-corrector are raised to a

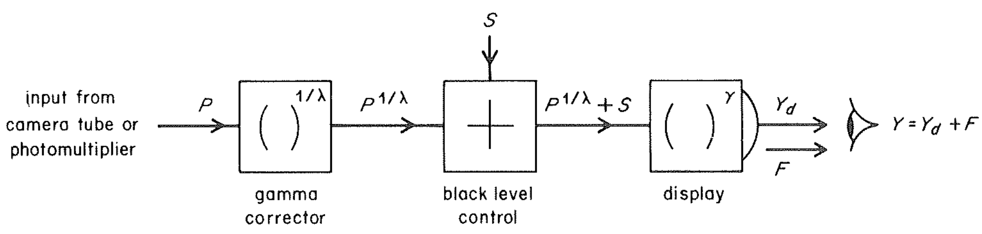

Fig. 5 - Signal chain — block diagram

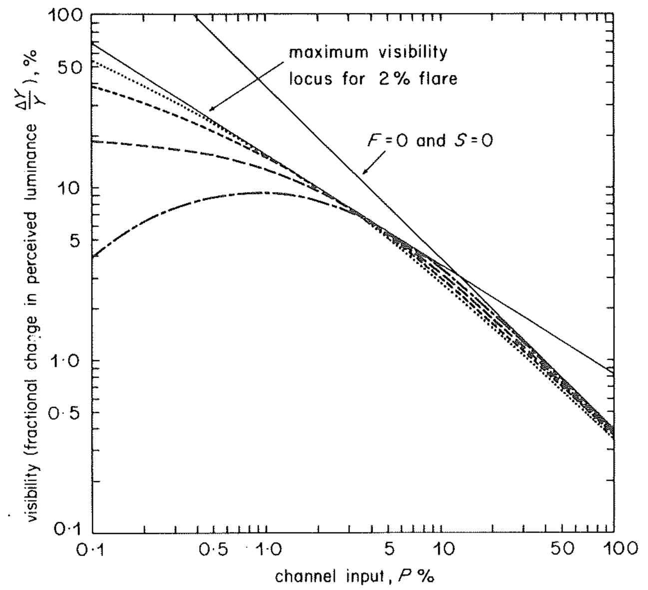

Fig. 6 - Visibility of 8 bit quantisation of channel input P

$F = 2\%$ { ············ $s = 0{\cdot}1$; – – – – – $s = 0{\cdot}05$; — — — $s = 0$; — - — $s = -0{\cdot}5$ }

display luminance where they are more visible. For each value of input P there is a worst case value of display black level (s) which makes quantisation at this input level most visible. It can be shown that this worst-case occurs when $(P^{1/\lambda} + s)^{\gamma} = F(\gamma - 1)$.

$$\text{Thus}\quad \frac{\Delta Y}{Y}_{max,\ s} = \frac{1}{2^n}\cdot\frac{(\gamma-1)^{(1-1/\gamma)}}{\lambda}\cdot\frac{P_{max}}{P^{(1-1/\lambda)}}\cdot\frac{1}{F^{1/\gamma}}$$

Thus the worst-case visibility of an input step depends on the level of input (P) at which the step occurs, as shown by the worst case visibility locus (Fig. 6). The worst-case visibility is greatest in the darkest parts of the picture (i.e. when $P = P_{min}$).

$$\text{Then}\quad \left(\frac{\Delta Y}{Y}\right)_{max\ s,\ p} = \frac{1}{2^n}\cdot\frac{(\gamma-1)^{(1-1/\gamma)}}{\lambda}\cdot\left(\frac{P_{max}}{P_{min}}\right)^{(1-1/\lambda)}\cdot\left(\frac{P_{max}^{\ \gamma/\lambda}}{F}\right)^{1/\gamma}$$

It is useful to compare this maximum visibility of input quantisation (with worst-case display black level) with the maximum visibility of gamma-corrected quantisation (which is independent of display black level) since more is at present known about the visibility of quantisation of gamma-corrected signals.

$$\frac{\left(\frac{\Delta Y}{Y}\right)^{input}_{max\ s,\ p}}{\left(\frac{\Delta Y}{Y}\right)^{gamma\ corrected}_{max}} = \frac{\frac{1}{2^n}\,\frac{(\gamma-1)^{(1-1/\gamma)}}{\lambda}\left(\frac{P_{max}}{P_{min}}\right)^{1-1/\lambda}\left(\frac{P_{max}^{\ \gamma/\lambda}}{F}\right)^{1/\gamma}}{\frac{(\gamma-1)^{(1-1/\gamma)}}{2^n}\cdot\left(\frac{Y_{d\ max}}{F}\right)^{1/\gamma}}$$

since $Y_{d\ max} \simeq P_{max}^{\ \gamma/\lambda}$ this reduces to

$$\frac{\left(\frac{\Delta Y}{Y}\right)^{input}_{max\ s,\ p}}{\left(\frac{\Delta Y}{Y}\right)^{gamma\ corrected}_{max}} = \frac{1}{\lambda}\left(\frac{P_{max}}{P_{min}}\right)^{(1-1/\lambda)}$$

e.g. if $\lambda = 2{\cdot}8$ and the input contrast ratio $(P_{max})/(P_{min})$ is 1000:1, then quantisation of the input would be 30 times more visible than quantisation of the gamma-corrected signals (or, for the same visibility, 5 more bits would be required for quantisation of the input signals). If it is assumed that 8 bits (or 5 bits plus dither) are necessary for satisfactory coding of gamma-corrected signals, then 13 bits (or 10 bits plus dither) would be required to code input signals over the full 1000:1 range if the display brightness were adjusted to worst-case.

2.6. Multiple quantisation

Figs. 1, 2, 3, 4 and 6 indicate the relative visibilities of the contours produced if one quantising process is performed at a number of alternative points in a processing channel. However, in practice, the digital signals produced after quantisation may be subject to further errors due to the subsequent processing which may result in quantising steps effectively larger than those produced by the primary quantisation process considered separately.

For example, the smooth curve (a) of Fig. 7 shows the ideal transfer characteristic of a simple gamma-corrector. If the input to the corrector is restricted to 16 discrete levels, as shown, the output will follow the dotted line, Fig. 7, curve (b), since the output will only change when the input moves from one discrete level to the next.

If the output of the corrector is also restricted to 16 discrete levels the output will be rounded up or down to one of these levels as shown by the solid line in Fig. 7, curve (c), this rounding of the output restricts the size of the output steps to a whole number of the 16 possible output levels. For example, when the input is greater than 4, the steps shown dotted (i.e. those due to input quantisation alone) are smaller than one output level and they are rounded either to zero (e.g. the input transitions 7–8, 10–11, 12–13, 15–16 produce no change in the output) or to one

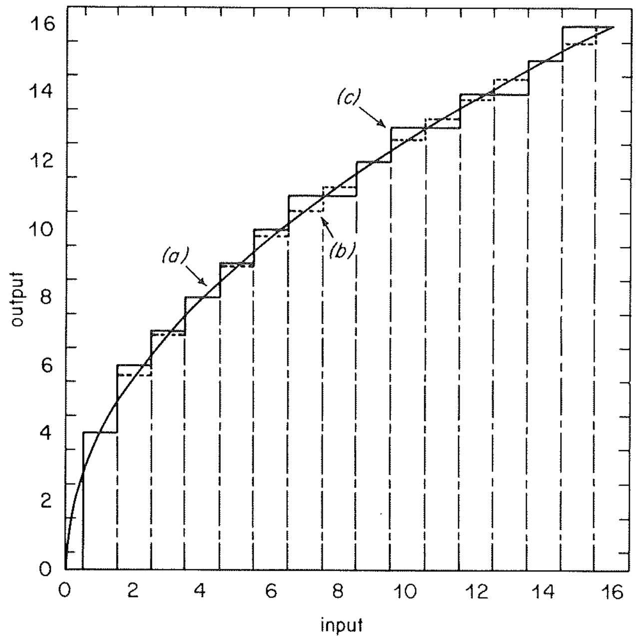

Fig. 7 - Transfer characteristic of gamma-corrector
(*a*) Ideal (*b*) Input quantised (*c*) Input and output quantised

complete output level (e.g. input transitions 4–5, 5–6, 6–7, 8–9, 9–10, 11–12, 13–14, 14–15 produce the output changes 8–9, 9–10, 10–11, 11–12, 12–13, 13–14, 14–15, 15–16 respectively). When the input is between 1 and 4 the steps shown dotted lie between 1 and 2 output levels. These may be rounded down to exactly one output step (e.g. input transitions 2–3 and 3–4 produce output changes 6–7 and 7–8), or may be increased to 2 output steps (e.g. the input transition 1–2 at the input).

Since the output-step size is always rounded up or down to a whole number of output steps, the visibility of

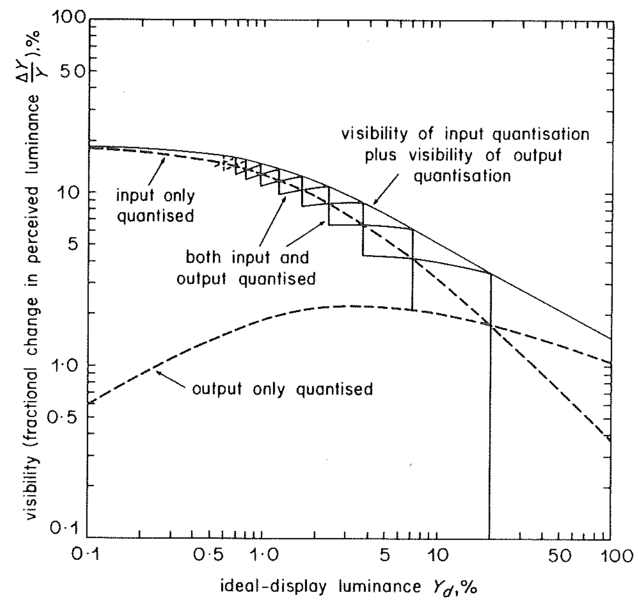

Fig. 8 - Visibility of quantisation of input and output of gamma-corrector

quantising both input and output must be an integer multiple of the visibility of quantising the output alone. The visibility of quantising the input must be rounded up or down to a multiple of the visibility of output steps.

This effect is illustrated in Fig. 8 for a gamma-corrector using 8 bits at both input and output. When Y_d is greater than 24% of $Y_{d\ max}$ quantisation of the gamma-corrected output is most visible and any step on the input signal is either increased to this size or ignored completely. When Y_d is less than 24% of $Y_{d\ max}$ the changes due to quantisation of the input signals are rounded to a whole number of output steps so that the resulting visibility is always an integer multiple of the visibility of one output step. Note that two visibility curves are given since each input step may be either rounded up or rounded down to a whole number of output steps; it is very difficult to predict whether a particular transition at the input will be rounded up or down.

A curve representing the sum of the two individual visibilities passes through the peaks of the composite curve and may be regarded as a worst-case design curve. If the total number of steps is small it may be possible to adjust the rounding so that most of the steps lie on the jagged double line well away from this worst-case curve, e.g. in the example quoted it is the transition from 0–1 at the input giving 0–4 at the output which would be most visible and it has been arranged so that no rounding takes place on this step. However, if a large number of bits is being used (as would be the case in any broadcast system) then this is not possible and it must be assumed that several of the transitions will lie on, or close to, the design curve (i.e. will have visibilities which are simply the sums of the visibilities of the two separate quantisations.

Thus the worst-case visibility of quantisations in a digital processor involving quantisation at several different points may be obtained by adding the visibilities of quantisation in each case. (Note the boundary between two regions of quantisation is marked by any digital process involving rounding or including multiplication, since this usually generates extra bits that are discarded.)

3. Theory-colour

So far consideration has been limited to the visibility of luminance steps, produced by quantising monochrome signals.

In a three colour system the displayed luminance Y_d is equal to a weighted average of the Red, Green and Blue signals: $Y_d = lR + mG + nB$. The Red, Green and Blue signals may be separately quantised each using n bits so that the difference between levels are:

$$\Delta R = \frac{R_{max}}{2^n},\ \Delta G = \frac{G_{max}}{2^n},\ \Delta B = \frac{B_{max}}{2^n}$$

If a transition occurs simultaneously in the Red, Green and Blue channels, the luminance change will be

$$\Delta Y_d = l\Delta R + m\Delta G + n\Delta B$$

$$= \frac{1}{2^n} (l\, R_{max} + m\, G_{max} + n\, B_{max})$$

$$\frac{\Delta Y_d}{Y} = \frac{1}{2^n} \cdot \frac{Y_{max}}{Y}$$ as before for the monochrome case.

This is the worst case which occurs only rarely because in general, transitions will not be exactly superimposed. This is especially true if dither is used to break up coherent transitions, the quantising errors in each channel then appear as uncorrelated noise. The luminance noise may be calculated by adding weighted noise powers from the Red, Green and Blue channels.

$$(\Delta Y)^2 = (l\Delta R)^2 + (m\Delta G)^2 + (n\Delta B)^2$$

if $\Delta R = \Delta G = \Delta B = \frac{R_{max}}{2^n} = \frac{G_{max}}{2^n} = \frac{B_{max}}{2^n} = \frac{Y_{max}}{2^n}$

$$\frac{\Delta Y}{Y} = \frac{1}{2^n} \frac{Y_{max}}{Y} (l^2 + m^2 + n^2)^{½}$$

$$= \frac{0{\cdot}6686}{2^n} \cdot \frac{Y_{max}}{Y} \quad (\text{if } l = 0{\cdot}299,\ m = 0{\cdot}587,\ n = 0{\cdot}114)$$

i.e. since the errors in the three channels are uncorrelated they will sometimes cancel to give less total luminance noise.

However, since the transitions do not occur at the same time there will be colour changes as well as luminance changes. The visibility of these colour changes may be assessed by considering their effect on the colour difference signals.

$$E_v = 0{\cdot}615R - 0{\cdot}515G - 0{\cdot}1B$$

$$E_u = -\,0{\cdot}147R - 0{\cdot}289G + 0{\cdot}437B$$

The total chrominance perturbation caused by uncorrelated perturbations ΔR, ΔG, ΔB is

$$\Delta E_{v,u} = [(0{\cdot}615\Delta R)^2 + (0{\cdot}515\Delta G)^2 + (0{\cdot}1\Delta B)^2 + (0{\cdot}147\Delta R)^2 + (0{\cdot}289\Delta G)^2 + (0{\cdot}437\Delta B)^2]^{½}$$

If $\Delta R = \Delta G = \Delta B = \frac{R_{max}}{2^n} = \frac{G_{max}}{2^n} = \frac{B_{max}}{2^n} = \frac{Y_{max}}{2^n}$ as before

$$\Delta E_{v,u} = \frac{0{\cdot}9744}{2^n} \cdot Y_{max}$$

Chrominance changes are normally assumed to be 6 dB less visible than luminance changes and, therefore, the total visibility of separately quantising Red, Green and Blue is:

$$\left[\left(\frac{\Delta Y}{Y}\right)^2 + \left(\frac{\Delta E_{v,u}}{2Y}\right)^2\right]^{½} =$$

$$= \left[0{\cdot}6686^2 + \left(\frac{0{\cdot}9744}{2}\right)^2\right]^{½} \frac{1}{2^n} \cdot \frac{Y_{max}}{Y}$$

$$= \frac{0{\cdot}8272}{2^n} \cdot \frac{Y_{max}}{Y}$$

This is still slightly less visible than the monochrome case.

It is also possible to calculate the effects of transitions in only one channel at a time.

The visibility of a change ΔR in the Red channel

$$= \left[0{\cdot}299^2 + \left(\frac{0{\cdot}615}{2}\right)^2 + \left(\frac{0{\cdot}147}{2}\right)^2\right]^{½} \frac{1}{2^n} \cdot \frac{Y_{max}}{Y}$$

$$= \frac{0{\cdot}435}{2^n} \cdot \frac{Y_{max}}{Y}$$

Similarly for $G = \frac{0{\cdot}657}{2^n} \cdot \frac{Y_{max}}{Y}$

and $B = \frac{0{\cdot}25}{2^n} \cdot \frac{Y_{max}}{Y}$

Quantisation of the green signal should therefore be more visible than that of red and this in turn should be more visible than that of blue.

The foregoing applies to quantisation of 'linear' R, G, B signals so that $\Delta R = \Delta G = \Delta B$ = constant. If gamma-corrected or logarithmic signals are quantised ΔR, ΔG and ΔB will no longer be constant and will not in general be equal unless $R = G = B$ (i.e. Grey). In this special case the monochrome results would be modified by the same visibility factors as above but for other colours the problem becomes too complex to analyse fully. However, it is likely that, as for linear quantisation, the monochrome results represent the worst-case and should be used for design purposes.

4. Subjective tests on the quantisation of channel inputs

Section 2 showed that quantisation of input signals is most visible in the dark areas of the picture. Since the visibility depends on the viewing conditions, the following subjective tests were conducted to determine the number of bits required under the most critical viewing conditions likely to be encountered in practice.

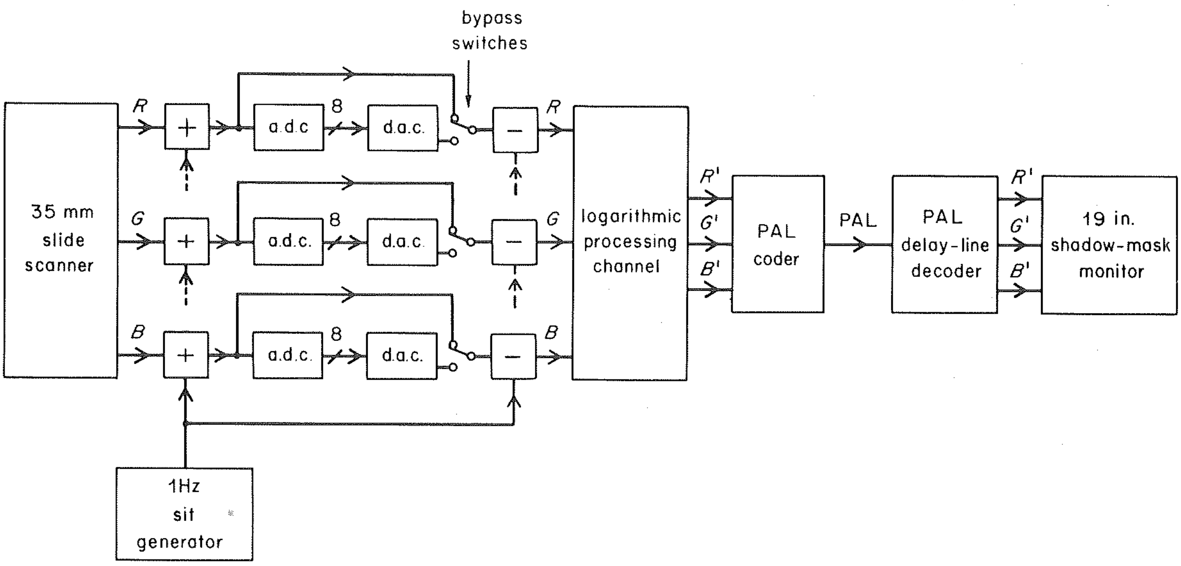

Fig. 9 - Test equipment — block diagram

4.1. Equipment

Three 8 bit ADC/DAC units[3] were placed in the Red, Green and Blue inputs of a logarithmic telecine processing channel as shown in Fig. 9.

After coding and decoding, using PAL System I, the outputs were displayed on a high quality 22 in. shadow-mask monitor. The lighting in the test room was adjusted to give an illumination (as measured by reflection from a magnesium carbonate block in front of the Monitor) of 1·6 lm/m^2, since this is the lowest value found in a survey[4] of domestic viewing conditions. For early tests the peak brightness of the monitor was set to 20 ft/lamberts (68 cd/m^2) as recommended by the CCIR[5] but this was increased in later tests to 50 ft/lamberts (171 cd/m^2) since this made quantisation more visible without degrading the picture quality in any other way (although few domestic receivers can produce such bright pictures this test condition allows scope for improved designs).

The monitor cut-off was adjusted using PLUGE and in the first tests the channel output 'sit' was adjusted to 0% (i.e. so that with no input to the channel the monitor was just cut-off). In later tests the channel output sit was increased to 10% since this made quantisation more visible (as theory predicts) without seriously degrading the pictures in any other way. Brief tests showed that quantisation was most visible with 17% sit but since the picture became unacceptable with more than 10% sit this was considered to be the worst-case that would be encountered in practice; 10% sit increased the visibility by 2 grades.

Care was taken to clamp the signals at the ADC input so that blanking did not contain any transitions which would produce streaking when reclamped by the processing channel. The ADC/DAC units were also modified slightly so that no overall d.c. shift was produced by removal of bits. Fig. 10 shows the half quantum-level mean d.c. shift normally produced by quantisation (this effect occurs

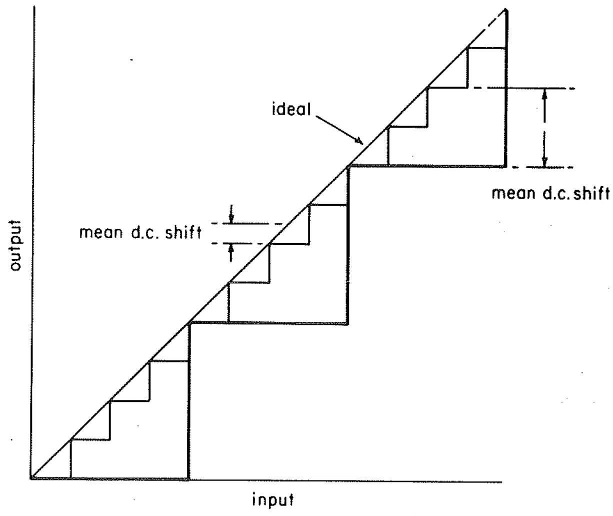

Fig. 10 - ADC/DAC transfer characteristic

because the number of bits is restricted by truncation rather than rounding). This effect, which produced some anomalies in the first tests, was overcome by 'sitting up' the input to the ADC by half a level to pre-correct for the truncation produced by the ADC. This was simply instrumented by adding blanking equal to half a quantum-level to the input of the ADC (so that the input was sat up by the ADC clamp).

4.2. Signal sources

The video input for the tests was obtained from a flying-spot slide scanner. The four colour slides used are reproduced in monochrome as Fig. 11(*a*), (*b*), (*c*) and (*d*). These slides were chosen because they all contain large areas of low luminance in which input quantisation is most visible. The input signal gain was adjusted to fill the ADC range.

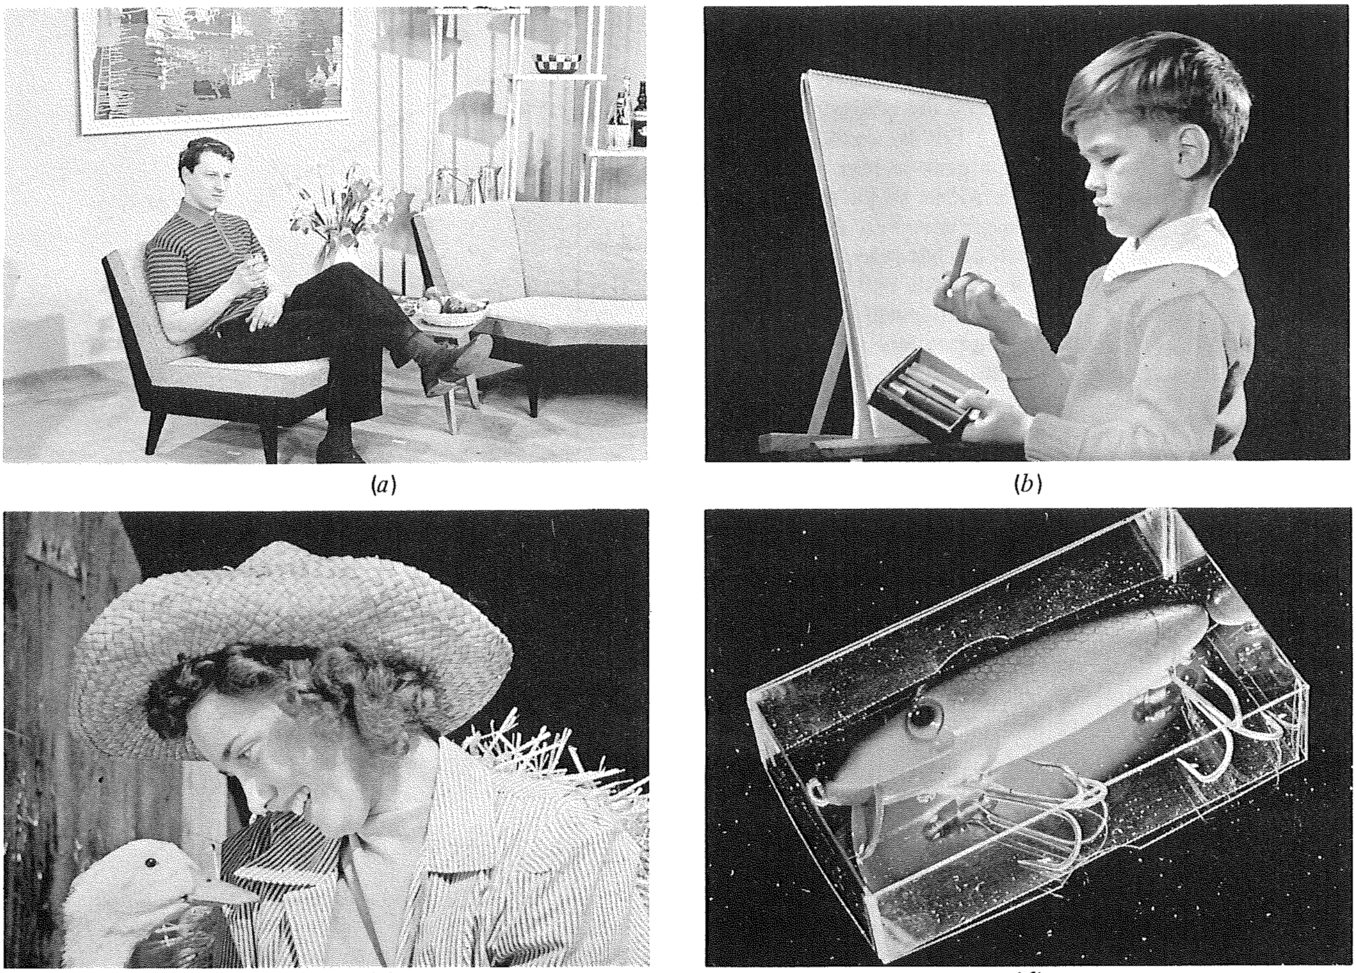

(a) (b) (c) (d)

Fig. 11 - Monochrome versions of slides used in tests
(*a*) Studio (*b*) Boy (*c*) Girl (*d*) Fish-hook

The use of stationary pictures made the tests easily repeatable and gave the observers time to find and examine the most critical parts of the picture. However, since quantising effects that move slowly are more visible than stationary ones, a simulation was produced by adding 1 Hz sinusoidally varying sit (of peak to peak magnitude = 1/256th peak white) to the scanner outputs. This made the level transitions move to and fro across the picture, thus making them more noticeable. Since in the early tests it was found that this 'periodic sit change' seriously impaired both quantised and unquantised pictures, in later tests the overall sit variation was cancelled after the ADC/DAC so that only the transitions appeared to move in an otherwise unvarying scene.

4.3. Test procedure

A total of 12 technical observers took part in the tests. They were seated at a distance of between five and seven times picture height from the display.

Before each session (which consisted of 30 tests occupying about 25 minutes) the effects of quantisation were demonstrated using the most visible test condition that was to occur in that session. The observers were then asked to grade each test using the EBU 6 point impairment scale given below:

Grade	Degree of Impairment
1	Imperceptible
2	Just perceptible
3	Definitely perceptible but not disturbing
4	Somewhat objectionable
5	Definitely objectionable
6	Unusable

Between each test the unquantised picture was shown for comparison and the observers were asked to ignore any impairment that also appeared on this reference.

5. Results of subjective tests

8 bit quantisation was found to give objectionable (Grade 4¼) pictures so dither was used to conceal quantisation.

5.1. The effectiveness of dither

The highly visible coarse patterns (contours) produced by quantisation can be broken up by the addition of dither waveforms to the ADC input. For example, if a half sampling-frequency signal with a peak-to-peak magnitude equal to half a quantum level is added to the ADC input, the output alternates between levels whenever the input lies close (±¼ quantum) to a transition between levels. The mean level of this alternating output lies half way between levels and is seen as a set of 'pseudo-levels' equivalent to one extra bit. (The half sampling-frequency component is removed by the low-pass filter following the DAC.)

Similarly, if a one-quarter sampling-frequency signal with a peak-to-peak magnitude equal to one quarter of one quantum level is also added to the input, the output alternates at one-quarter sampling-frequency between levels and pseudo-levels whenever the input signal level lies close (± one eighth of one quantum level) to a pseudo-transition. This creates a second set of pseudo-levels and gives the impression that another extra bit has been used. (The one-quarter sampling-frequency component is not removed by filtering but being at high frequency is less visible than the coarse patterns which it conceals.)

The remaining contouring can be further broken up by adding random noise to the input (a certain amount of noise may in any case be present on the input signal). This causes the output to fluctuate between levels (and pseudo-levels) in a random manner which smoothes out the remaining contours but leaves a certain amount of noise on the picture. Thus by adding dither to the ADC input, coherent (picture dependent) contouring can be transformed into h.f. patterning and incoherent noise.

The three forms of dither described above were added to 6, 7 and 8 bit quantisation of all three linear colour-separation signals. The results of 56 assessments are shown in Fig. 12. It will be seen that when added to 6 and 7 bit quantisation, half sampling-frequency (½ f.s.) dither is equivalent to slightly more than one extra bit; the reason for this bonus is not known. Quarter sampling-frequency (¼ f.s.) dither is equivalent to a further two-thirds of a bit and noise to a further half a bit. Thus all three dithers are together equivalent to an extra 2½ bits (≙ 15 dB).

The beneficial effect of noise was less on 7 bits and slightly negative on 8 bits. It is thought that this was due to the presence of noise on the input signal, making extra noise unnecessary and indeed, for 8 bits, undesirable.

5.2. The relative visibility of Red, Green and Blue quantisation

The following tests were performed to check the conclusions of Section 3 (that quantisation of the Green signal is more visible than that of Red, which in turn is more visible than that of Blue).

Two of the colour-separation signals were quantised using 8 bits plus all three dithers, while the remaining signal

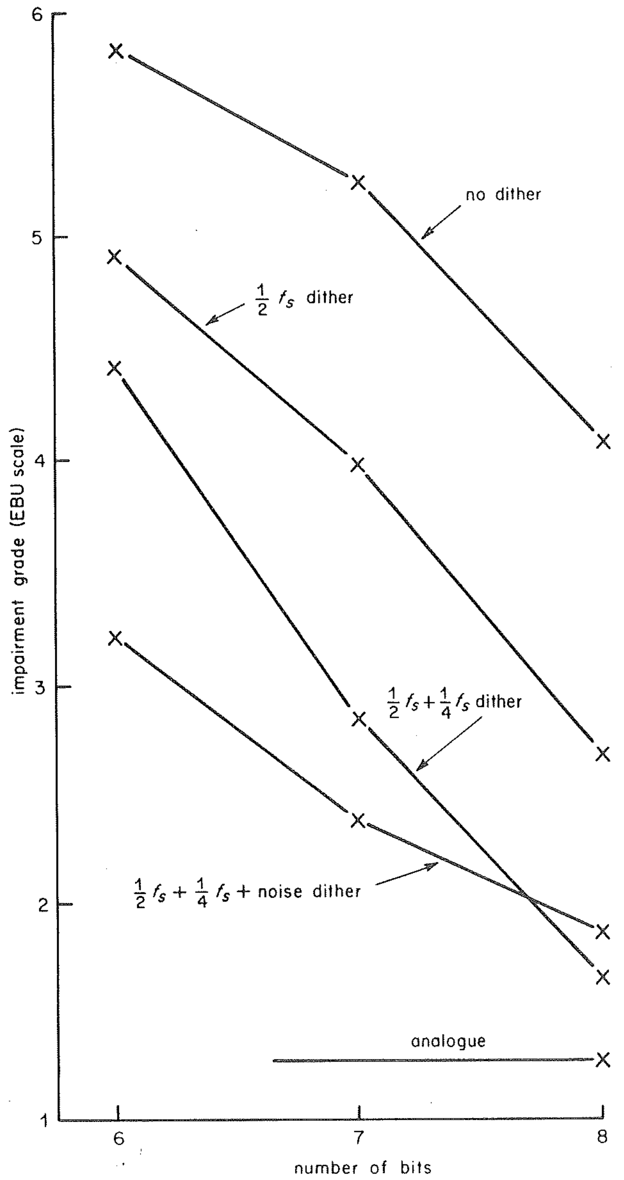

Fig. 12 - The effects of dither

was quantised to various degrees of accuracy. These tests were carried out using all four slides and then repeated using the most critical two slides as controls in an investigation into the effect of masking (see next Section). The averaged results of all these tests (in which masking was not used) is given in Fig. 13.

As expected from the theory, quantisation of the green signal was much more visible (by approximately 1½ bits) than that of blue. Quantisation of the red signal however was only slightly more visible than that of blue, whereas theory (Section 3) predicts a difference of three-quarters of a bit. The red input signal was, however, slightly more noisy than the green and blue and this extra noise acting as dither could, in part, explain the reduced visibility. A slight difference in the relative visibilities of ΔE_v and ΔE_u (assumed equal in Section 3) could also be responsible.

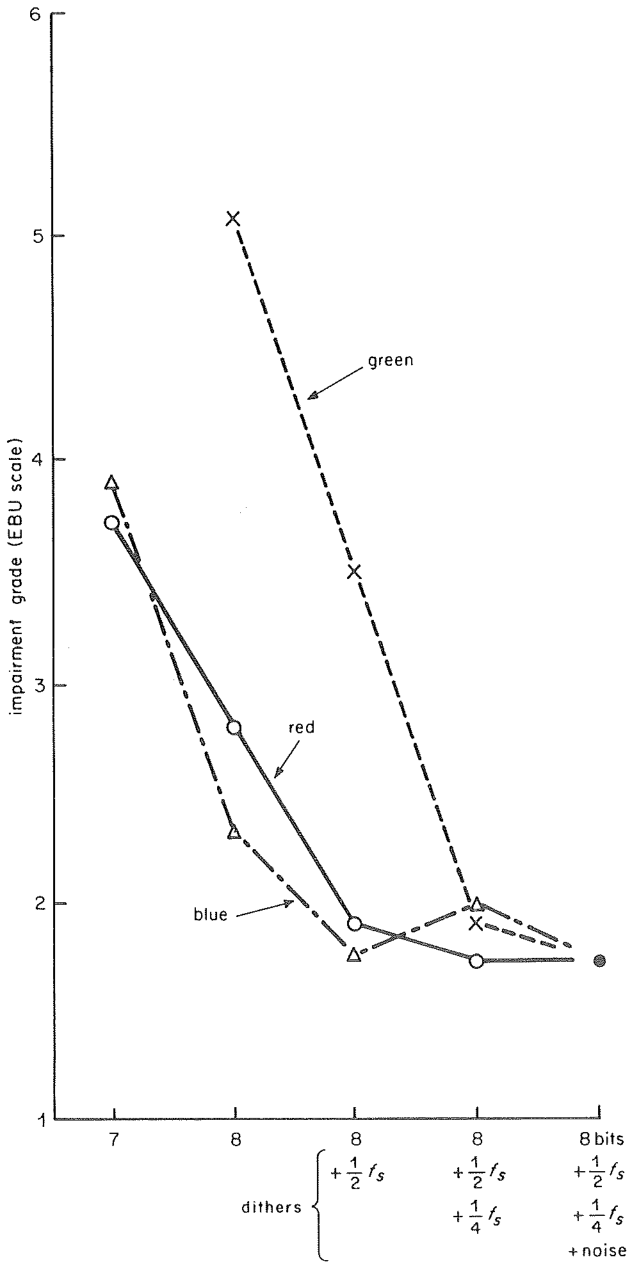

Fig. 13 - The relative visibilities of red, green and blue quantisations

5.3. The effect of masking

In a telecine processing channel the three colour-separation signals are matrixed together (in logarithmic form) in order to improve the overall colorimetry. This process (called masking) normally enhances colour differences and would be expected to make quantisation of the channel inputs more visible.

A series of tests, similar to those described in the last Section, was conducted with the following typical masking matrix in circuit.

$$\begin{vmatrix} \log R_{out} \\ \log G_{out} \\ \log B_{out} \end{vmatrix} = \begin{vmatrix} 1{\cdot}202 & -0{\cdot}053 & -0{\cdot}110 \\ -0{\cdot}206 & 1{\cdot}416 & -0{\cdot}217 \\ -0{\cdot}061 & -0{\cdot}549 & 1{\cdot}712 \end{vmatrix} \times \begin{vmatrix} \log R_{in} \\ \log G_{in} \\ \log B_{in} \end{vmatrix}$$

In order that the effect of the mask could be assessed the tests were interspersed (in groups of 3) with a similar set of 'control' tests in which masking was not employed. As before the subjects were asked to compare the quantised signal with the appropriate (i.e. masked or non-masked) unquantised reference which was shown between tests. Each set of tests was in random order but individual test pairs were arranged to be fairly close together to minimise the effects of learning and fatigue.

The two most critical slides as determined by previous tests ((*b*) Boy and (*d*) Fish Hook) were used for these tests and the results are given in Fig. 14. As expected, masking tended to produce an increase in the visibilities of the effects of quantising for all three colour-separation signals.

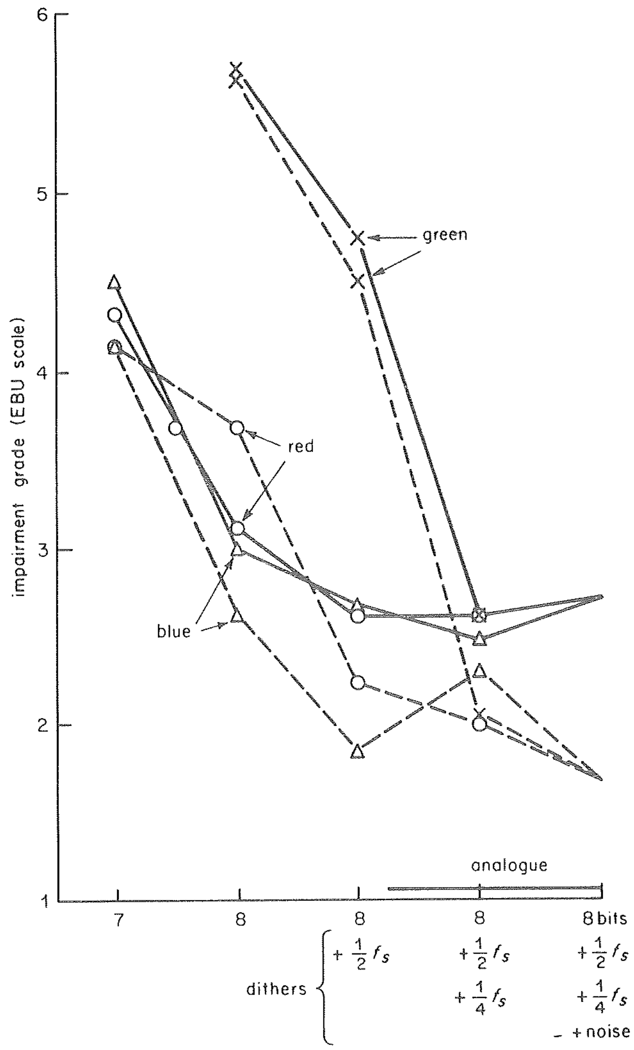

Fig. 14 - The effects of masking
(*a*) ——— masking (*b*) – – – – no masking

6. Conclusions

Quantisation of gamma-corrected signals is known from earlier work to be most visible in the mid-greys and this is confirmed by the theoretical analysis of Section 2.3.. This analysis also predicts that quantisation of logarithmic signals will be most visible at peak white and that this worst-case visibility will roughly equal that of gamma-corrected quantisation (i.e. requiring a minimum of 5 bits plus dither). This has been experimentally verified.

The theoretical analysis predicted that quantisation of non-gamma corrected (channel input) signals would be extremely visible in the black and that the exact visibility would depend on the viewing conditions and display adjustments. This proved to be the case.

Under worst-case conditions the best that could be achieved with the available 8 bit ADCs, using all available dither waveforms and with the input gain adjusted to fill the ADC range, was grade 2·1 without masking and 2·8 with masking. By extrapolation an extra bit would be required in the green channel to make the quantised signal indistinguishable from the analogue signal under these conditions. However, since the tests were conducted under the most rigorous conditions it will probably be possible to derive just-acceptable pictures under normal conditions using 8 bits with all available dither waveforms. In practice, 10 bits would provide a safety margin.

These figures make no allowance for the effects of subsequent quantisation since this would most likely be performed on logarithmic or gamma-corrected signals and therefore would have little effect on the visibility of quantisation in the dark areas of the picture. It should be noted that the above were conducted with the input signal gain adjusted to fill the ADC range. If this were not done the ADCs would not operate efficiently and quantisation would be more visible necessitating the use of ADCs with perhaps 11 bit resolution.

7. References

1. DEVEREUX, V.G. 1971. Pulse code modulation of video signals: subjective study of coding parameters. BBC Research Department Report No. 1971/40.

2. GEDDES, W.K.E. 1962. The variation in visibility of interference over the grey scale of a television picture. BBC Research Department Report No. T-092, Serial No. 1962/27.

3. DEVEREUX, V.G. 1970. Pulse code modulation of video signals: 8 bit coder and decoder. BBC Research Department Report No. 1970/25.

4. ROBINSON, R.N. 1972. The contrast handling requirements of a colour television display. BBC Research Department Report No. 1972/37.

5. CCIR 1972. Method for subjective assessment of the quality of television pictures. Conclusions of the Interim Meeting of Study Group 11 (Broadcasting service (Television) Geneva 5–18 July 1972. Recommendation AA/11.

SMW/JUC

Printed by BBC RESEARCH DEPARTMENT, Kingswood Warren, Tadworth, Surrey, KT20 6NP